



US010043078B2

(12) **United States Patent**
Winter et al.

(10) **Patent No.:** **US 10,043,078 B2**
(45) **Date of Patent:** **Aug. 7, 2018**

(54) **VIRTUAL TURNSTILE SYSTEM AND METHOD**

(71) Applicant: **Placemeter Inc.**, New York, NY (US)

(72) Inventors: **Alexandre Winter**, New York, NY (US); **Ugo Jardonnet**, New York, NY (US); **Tuan Hue Thi**, New York, NY (US); **Niklaus Haldimann**, New York, NY (US)

(73) Assignee: **Placemeter LLC**, New York, NY (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **15/078,611**

(22) Filed: **Mar. 23, 2016**

(65) **Prior Publication Data**

US 2016/0314353 A1 Oct. 27, 2016

Related U.S. Application Data

(63) Continuation-in-part of application No. 14/727,321, filed on Jun. 1, 2015.

(60) Provisional application No. 62/150,646, filed on Apr. 21, 2015.

(51) **Int. Cl.**

G06K 9/00 (2006.01)
G06K 9/62 (2006.01)
G06T 7/246 (2017.01)
G06T 7/194 (2017.01)

(52) **U.S. Cl.**

CPC **G06K 9/00744** (2013.01); **G06K 9/00771** (2013.01); **G06K 9/6267** (2013.01); **G06T 7/194** (2017.01); **G06T 7/246** (2017.01); **G06T 2207/10016** (2013.01); **G06T 2207/30232** (2013.01); **G06T 2207/30241** (2013.01); **G06T 2207/30242** (2013.01)

(58) **Field of Classification Search**

CPC .. G06Q 10/101; G06F 17/246; G06K 9/6267; G06K 9/00744; G06K 9/00771; G06T 7/194; G06T 7/246

See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,194,908 A * 3/1993 Lougheed G01P 3/38 250/332

6,587,574 B1 7/2003 Jeannin
6,987,883 B2 1/2006 Lipton et al.
6,999,600 B2 2/2006 Venetianer et al.
7,199,798 B1 4/2007 Echigo

(Continued)

FOREIGN PATENT DOCUMENTS

CN 101489121 A 7/2009

OTHER PUBLICATIONS

Machado, D. "People Counting System using Existing Surveillance Video Camera" Nov. 2011, pp. 1-71.*

(Continued)

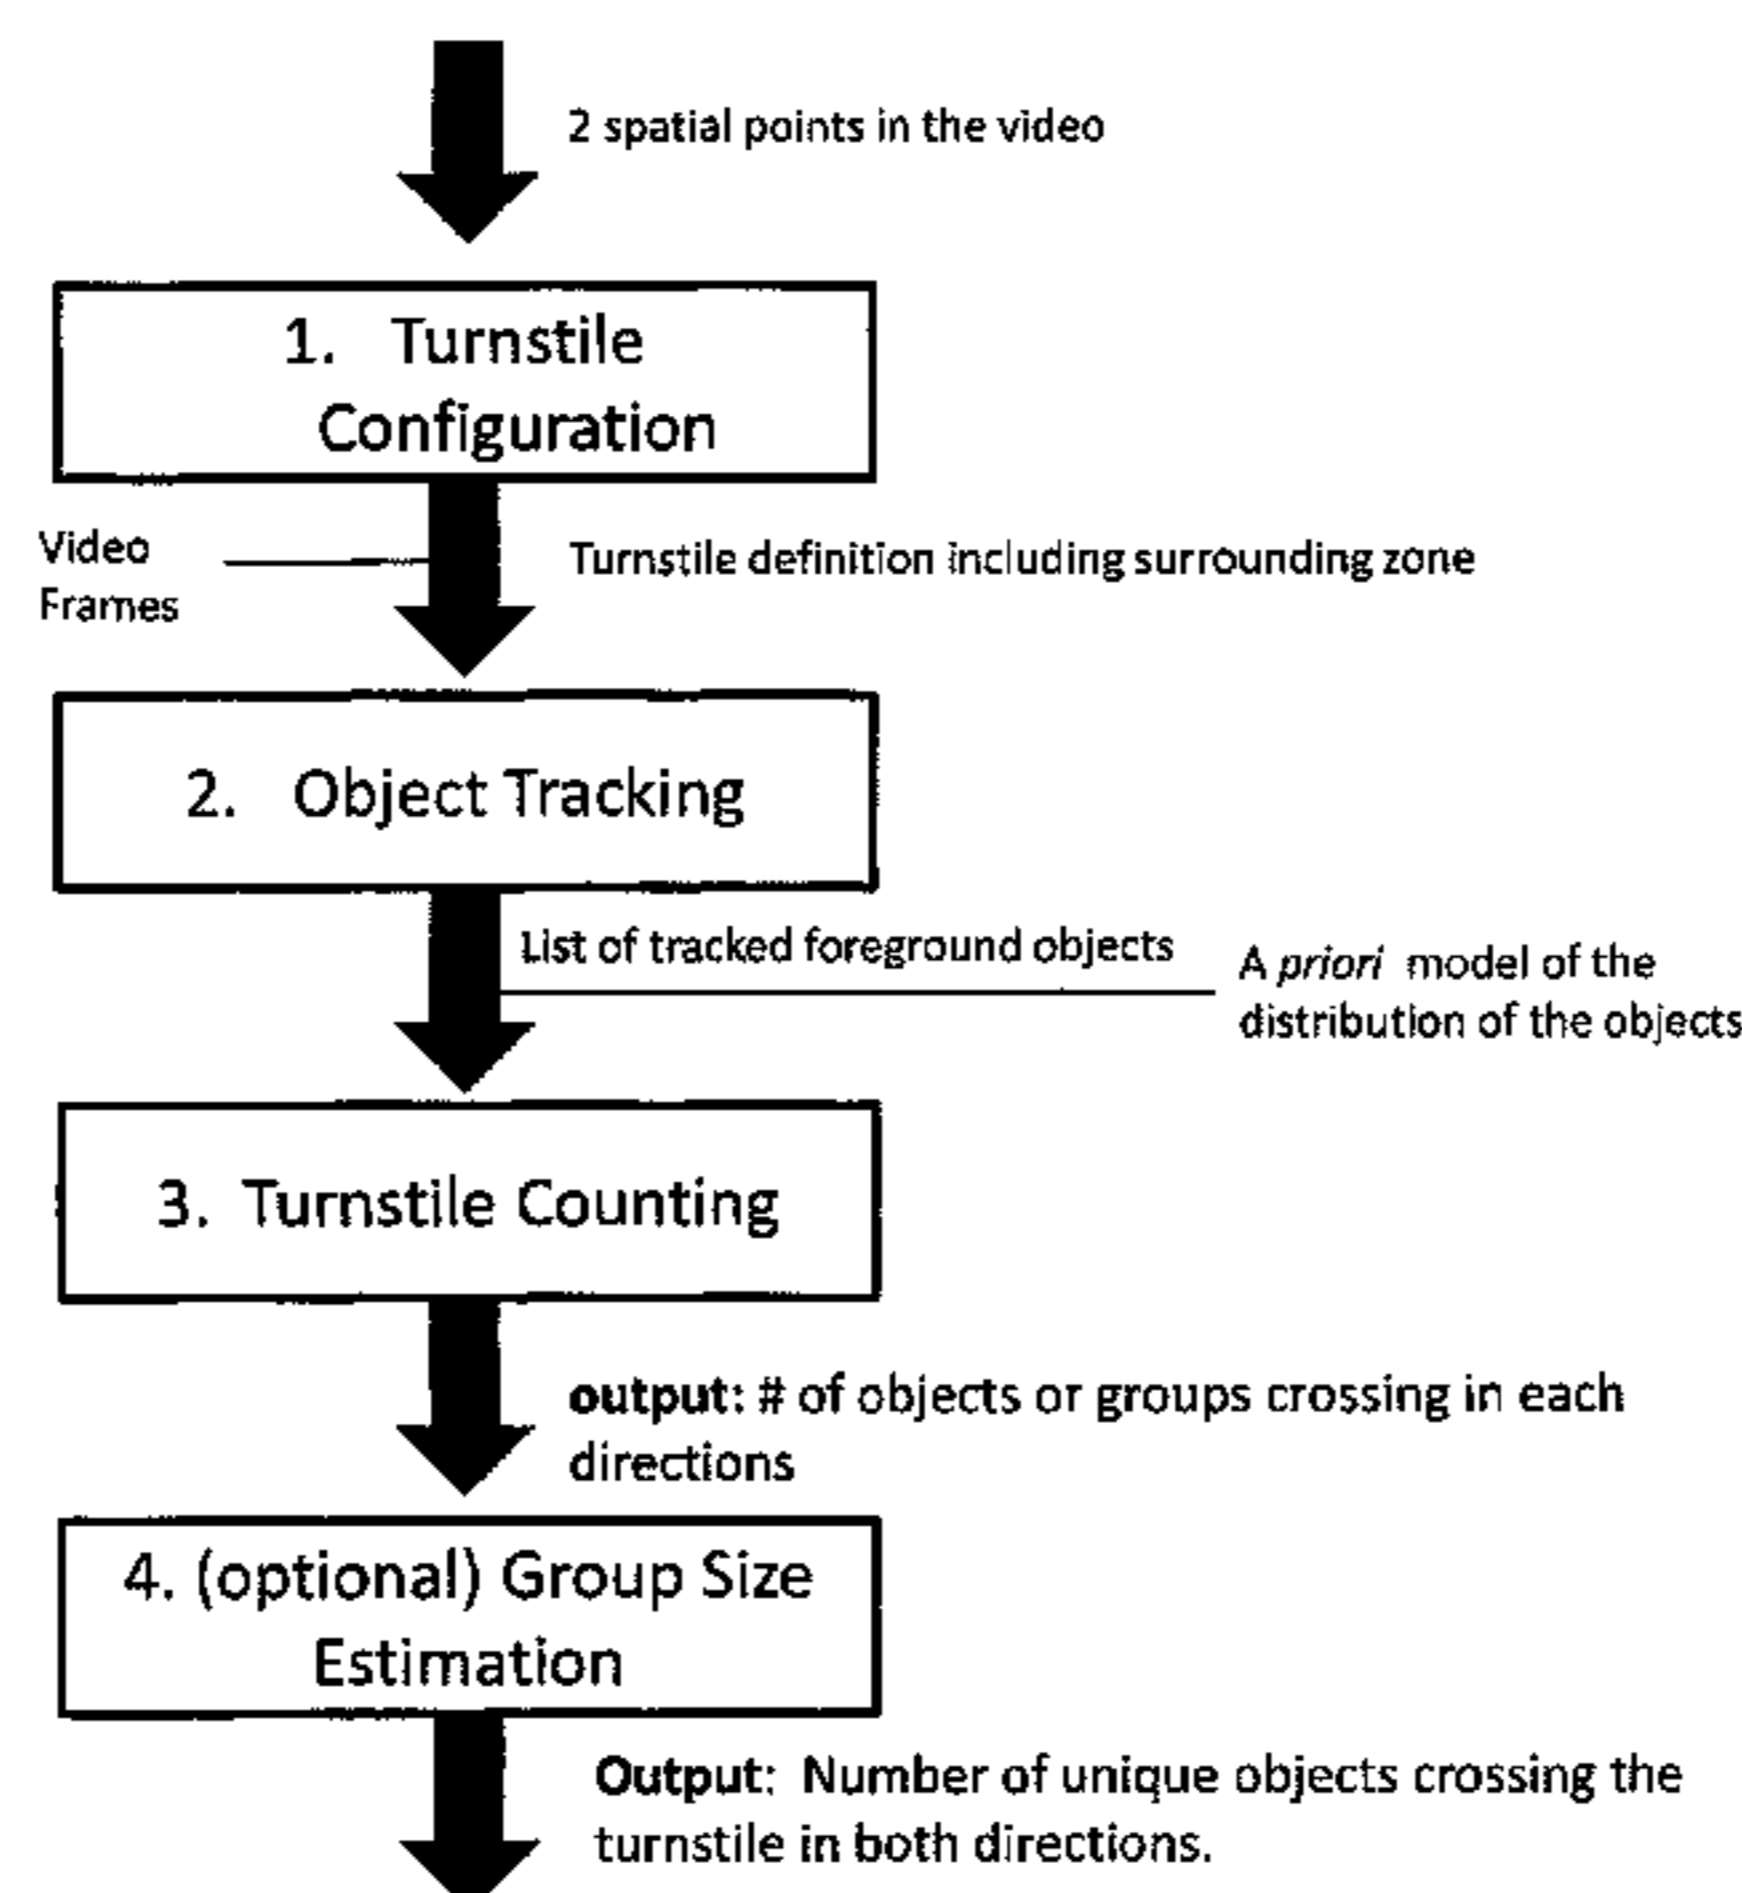
Primary Examiner — Mia M Thomas

(74) *Attorney, Agent, or Firm* — Barbara B Courtney; Courtney IP Law

(57) **ABSTRACT**

Embodiments include a system and method for activity monitoring using video data from multiple dissimilar sources. The video data is processed to remove any dependency of the system on types of video input data. The video data is processed to yield useful human readable information regarding events in real time, such as how many people move through a line in a period of time.

18 Claims, 5 Drawing Sheets



(56)

References Cited

U.S. PATENT DOCUMENTS

7,221,367 B2 5/2007 Cardno
 7,224,852 B2 5/2007 Lipton et al.
 7,391,907 B1 6/2008 Venetianer et al.
 7,424,175 B2 9/2008 Lipton et al.
 7,574,043 B2 8/2009 Porkili
 7,613,322 B2 11/2009 Yin et al.
 7,688,349 B2 3/2010 Flickner et al.
 7,796,780 B2 9/2010 Lipton et al.
 7,801,330 B2 9/2010 Zhang et al.
 7,868,912 B2 1/2011 Venetianer et al.
 7,932,923 B2 4/2011 Lipton et al.
 8,325,036 B1* 12/2012 Fuhr G06Q 10/087
 340/5.92
 8,331,619 B2 12/2012 Ikenoue
 8,340,349 B2 12/2012 Salgian
 8,340,654 B2 12/2012 Bratton et al.
 8,369,399 B2 2/2013 Egnal et al.
 8,401,229 B2 3/2013 Hassan-Shafique et al.
 8,457,401 B2 6/2013 Lipton et al.
 8,526,678 B2 9/2013 Liu et al.
 8,564,661 B2 10/2013 Lipton et al.
 8,582,803 B2 11/2013 Ding
 8,594,482 B2* 11/2013 Fan G08B 13/196
 386/200
 8,599,266 B2 12/2013 Trivedi
 8,625,905 B2 1/2014 Schmidt
 8,649,594 B1 2/2014 Hua
 8,654,197 B2 2/2014 Nizko
 8,655,016 B2 2/2014 Brown
 8,711,217 B2 4/2014 Venetianer et al.
 8,823,804 B2 9/2014 Haering et al.
 8,948,458 B2 2/2015 Hassan-Shafique et al.
 2003/0215110 A1* 11/2003 Rhoads G01C 11/00
 382/100
 2004/0022227 A1* 2/2004 Lynch G06Q 10/06
 370/338
 2004/0151342 A1* 8/2004 Venetianer G06K 9/00624
 382/103
 2004/0161133 A1* 8/2004 Elazar G01S 3/7864
 382/115
 2005/0169531 A1 8/2005 Fan
 2005/0185058 A1* 8/2005 Sablak H04N 5/23248
 348/208.99
 2006/0007308 A1* 1/2006 Ide G01S 3/7864
 348/143
 2006/0031291 A1* 2/2006 Beckemeyer G06Q 10/10
 709/204
 2006/0078047 A1* 4/2006 Shu G06F 17/3079
 375/240.01
 2006/0233535 A1* 10/2006 Honda G11B 27/034
 386/248
 2006/0290779 A1* 12/2006 Reverte E03F 7/10
 348/84
 2007/0024705 A1* 2/2007 Richter H04N 7/17318
 348/142
 2007/0024706 A1* 2/2007 Brannon, Jr. H04N 7/17318
 348/142
 2007/0071403 A1* 3/2007 Urita G11B 20/00086
 386/241
 2007/0016345 A1 5/2007 Peterson
 2007/0116345 A1* 5/2007 Peterson G06T 7/0081
 382/131
 2007/0127508 A1* 6/2007 Bahr H04N 7/17318
 370/401
 2007/0147690 A1 6/2007 Ishiwata
 2007/0177792 A1 8/2007 Ma
 2008/0030429 A1* 2/2008 Hailpern A63F 13/10
 345/8
 2008/0123955 A1* 5/2008 Yeh H04N 5/76
 382/173
 2008/0281518 A1* 11/2008 Dozier G07C 5/008
 701/469

2008/0316327 A1* 12/2008 Steinberg G06K 9/00248
 348/222.1
 2008/0316328 A1* 12/2008 Steinberg G06K 9/036
 348/222.1
 2009/0063205 A1 3/2009 Shibasaki
 2009/0103812 A1* 4/2009 Diggins H04N 19/61
 382/191
 2009/0141939 A1* 6/2009 Chambers G08B 13/19613
 382/103
 2009/0290023 A1* 11/2009 Lefort G08B 13/19608
 348/151
 2010/0177194 A1* 7/2010 Huang G06T 7/277
 348/157
 2010/0211304 A1 8/2010 Hwang
 2010/0290710 A1* 11/2010 Gagvani G06T 7/2053
 382/224
 2011/0130905 A1* 6/2011 Mayer G07C 5/008
 701/22
 2011/0141227 A1* 6/2011 Bigioi G06T 7/0075
 348/36
 2011/0152645 A1 6/2011 Hoover
 2011/0153645 A1* 6/2011 Hoover G06Q 30/02
 707/769
 2011/0280547 A1* 11/2011 Fan G08B 13/196
 386/239
 2012/0057640 A1* 3/2012 Shi H04N 19/52
 375/240.26
 2012/0086568 A1 4/2012 Scott
 2012/0127262 A1* 5/2012 Wu H04N 7/152
 348/14.09
 2012/0134535 A1* 5/2012 Pai G06K 9/00711
 382/103
 2012/0304805 A1 11/2012 Kim
 2013/0058537 A1* 3/2013 Chertok G06K 9/3233
 382/106
 2013/0086389 A1* 4/2013 Suwald G06F 21/34
 713/185
 2013/0148848 A1 6/2013 Lee
 2013/0202165 A1* 8/2013 Wehnes G06T 7/0012
 382/128
 2013/0251216 A1* 9/2013 Smowton H04L 9/3231
 382/118
 2014/0003708 A1* 1/2014 Datta G06K 9/6256
 382/159
 2014/0015846 A1* 1/2014 Campbell G06F 3/147
 345/589
 2014/0036090 A1* 2/2014 Black H04N 21/44016
 348/159
 2014/0278068 A1 9/2014 Stompolos
 2014/0294078 A1* 10/2014 Seregin H04N 19/00684
 375/240.15
 2014/0307056 A1* 10/2014 Collet Romea H04N 13/0007
 348/47
 2014/0359576 A1* 12/2014 Rath G06F 8/34
 717/121
 2015/0070506 A1* 3/2015 Chattopadhyay .. G06K 9/00718
 348/159
 2015/0077218 A1* 3/2015 Chakkaew G07C 9/00126
 340/5.2
 2015/0138332 A1* 5/2015 Cheng G06K 9/00771
 348/77
 2015/0339532 A1* 11/2015 Sharma A63F 13/213
 345/633
 2016/0012465 A1* 1/2016 Sharp G06Q 20/18
 705/14.17
 2016/0019427 A1* 1/2016 Martin G06K 9/00771
 382/103

OTHER PUBLICATIONS

Lefloch, D. "Real-Time People Counting system using Video Camera" Master of Computer Science, Image and Artificial Intelligence 2007 at UFR Sciences et Techniques, pp. 1-60.*
 Terada, K. "A method of counting the passing people by using the stereo images" Image Processing, 1999. ICIIP 99. Proceedings. 1999

(56)

References Cited

OTHER PUBLICATIONS

International Conference, pp. 1-5.*

Virtual Turnstile | Security Systems Technology | www.sstroup.co.uk/solutions/access-control/virtual-turnstile pp. 1-2.*

H. Bannour, L. Hlaoua and B. Ayeb, Survey of the Adequate Descriptor for Content-Based Image Retrieval on the Web: Global versus Local Features, Department of Information Sciences.

E. Baudrier, G. Millon, F. Nicolier and S. Ruan, A Fast Binary-Image Comparison Method with Local-dissimilarity Quantification, Laboratory Crestic, Troyes Cedex, France.

D. Lisin, M. Mattar, M. Blaschko, M. Benfield and E. Learned-Miller, Combining Local and Global Image Features for Object Class Recognition, Computer Vision Lab, Department of Computer Science, University of Massachusetts.

L. Paninski, Estimation of Entropy and Mutual Information, Center for Neural Science, New York University, New York, NY 10003, U.S.A., Received Jun. 3, 2002; accepted Nov. 27, 2002.

Machado, D. People Counting System using Existing Surveillance Video Camera: Nov. 2011, pp. 1-71.

Lefloch, D. "Real-Time People Counting System Using Video Camera" Master of Computer Science, Image and Artificial Intelligence 2007 at UFR Sciences et Technique, pp. 1-5.

Terada, K. "A method of counting the passing people by using the stereo images" Image Processing, 1999, ICIP 99. Proceedings. 1999 International Conference, pp. 1-5.

International Search Report and Written Opinion in PCT/US16/28511.

International Search Report and Written Opinion in PCT/US16/28516.

* cited by examiner

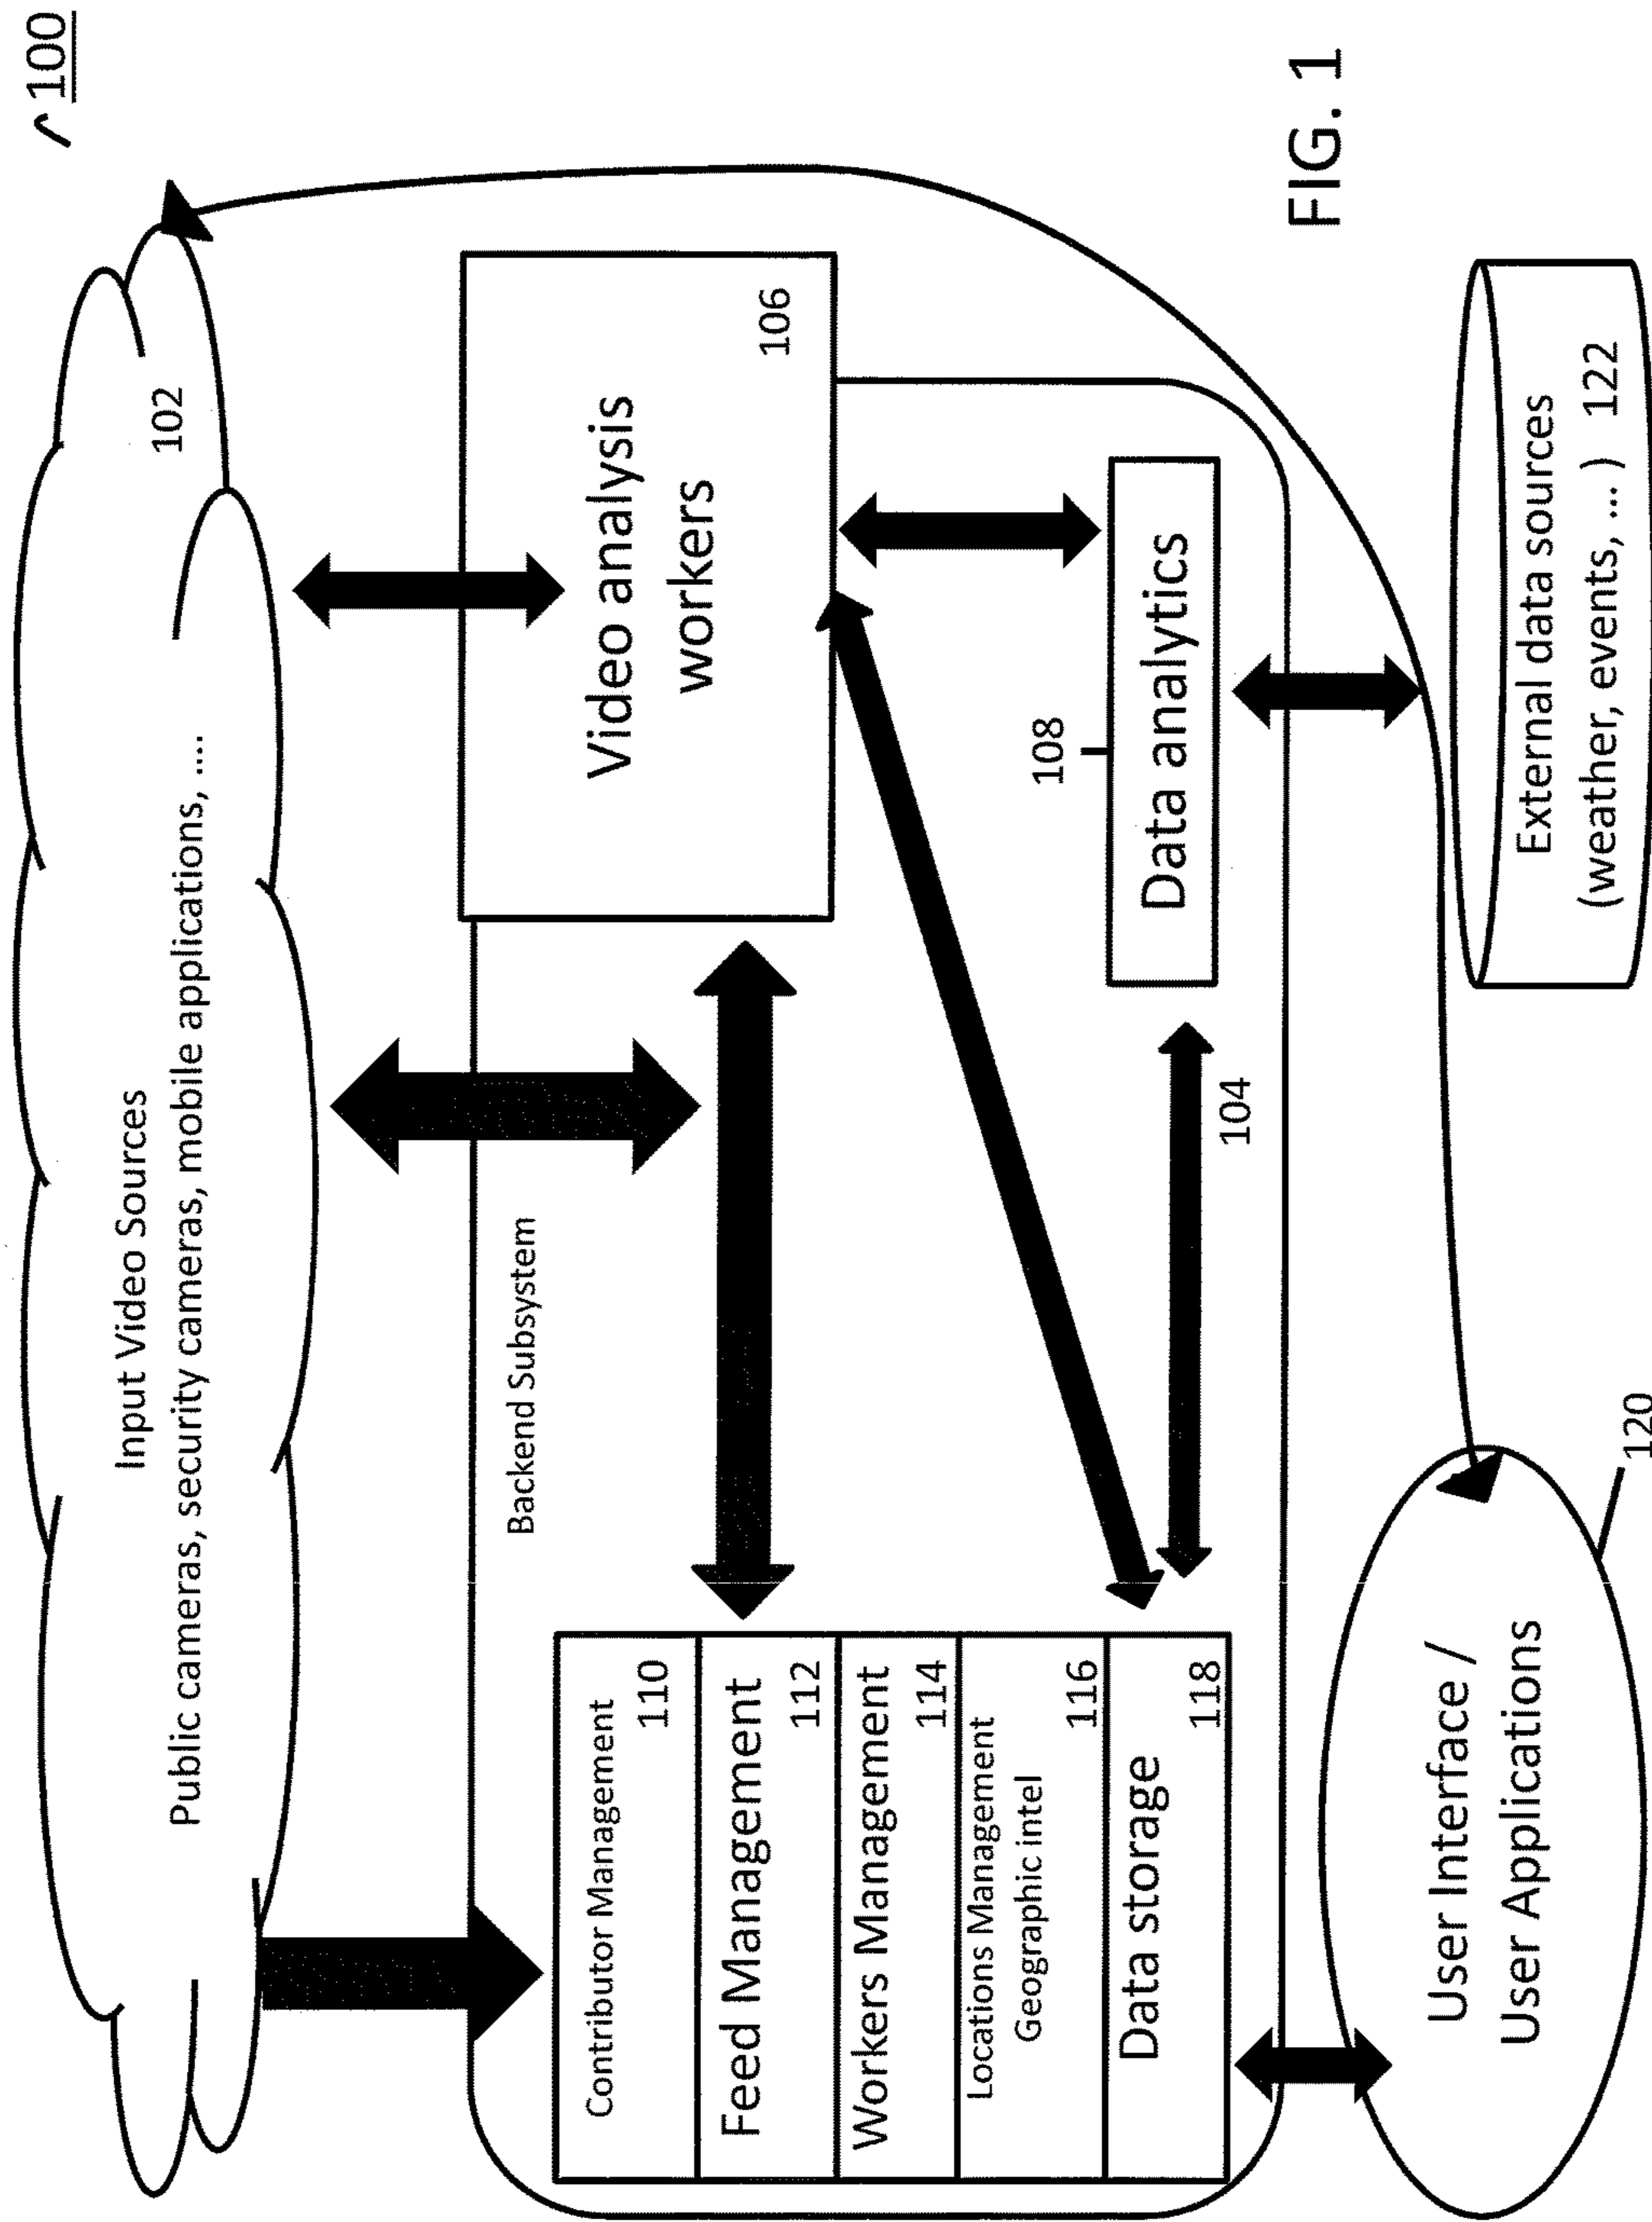


FIG. 1

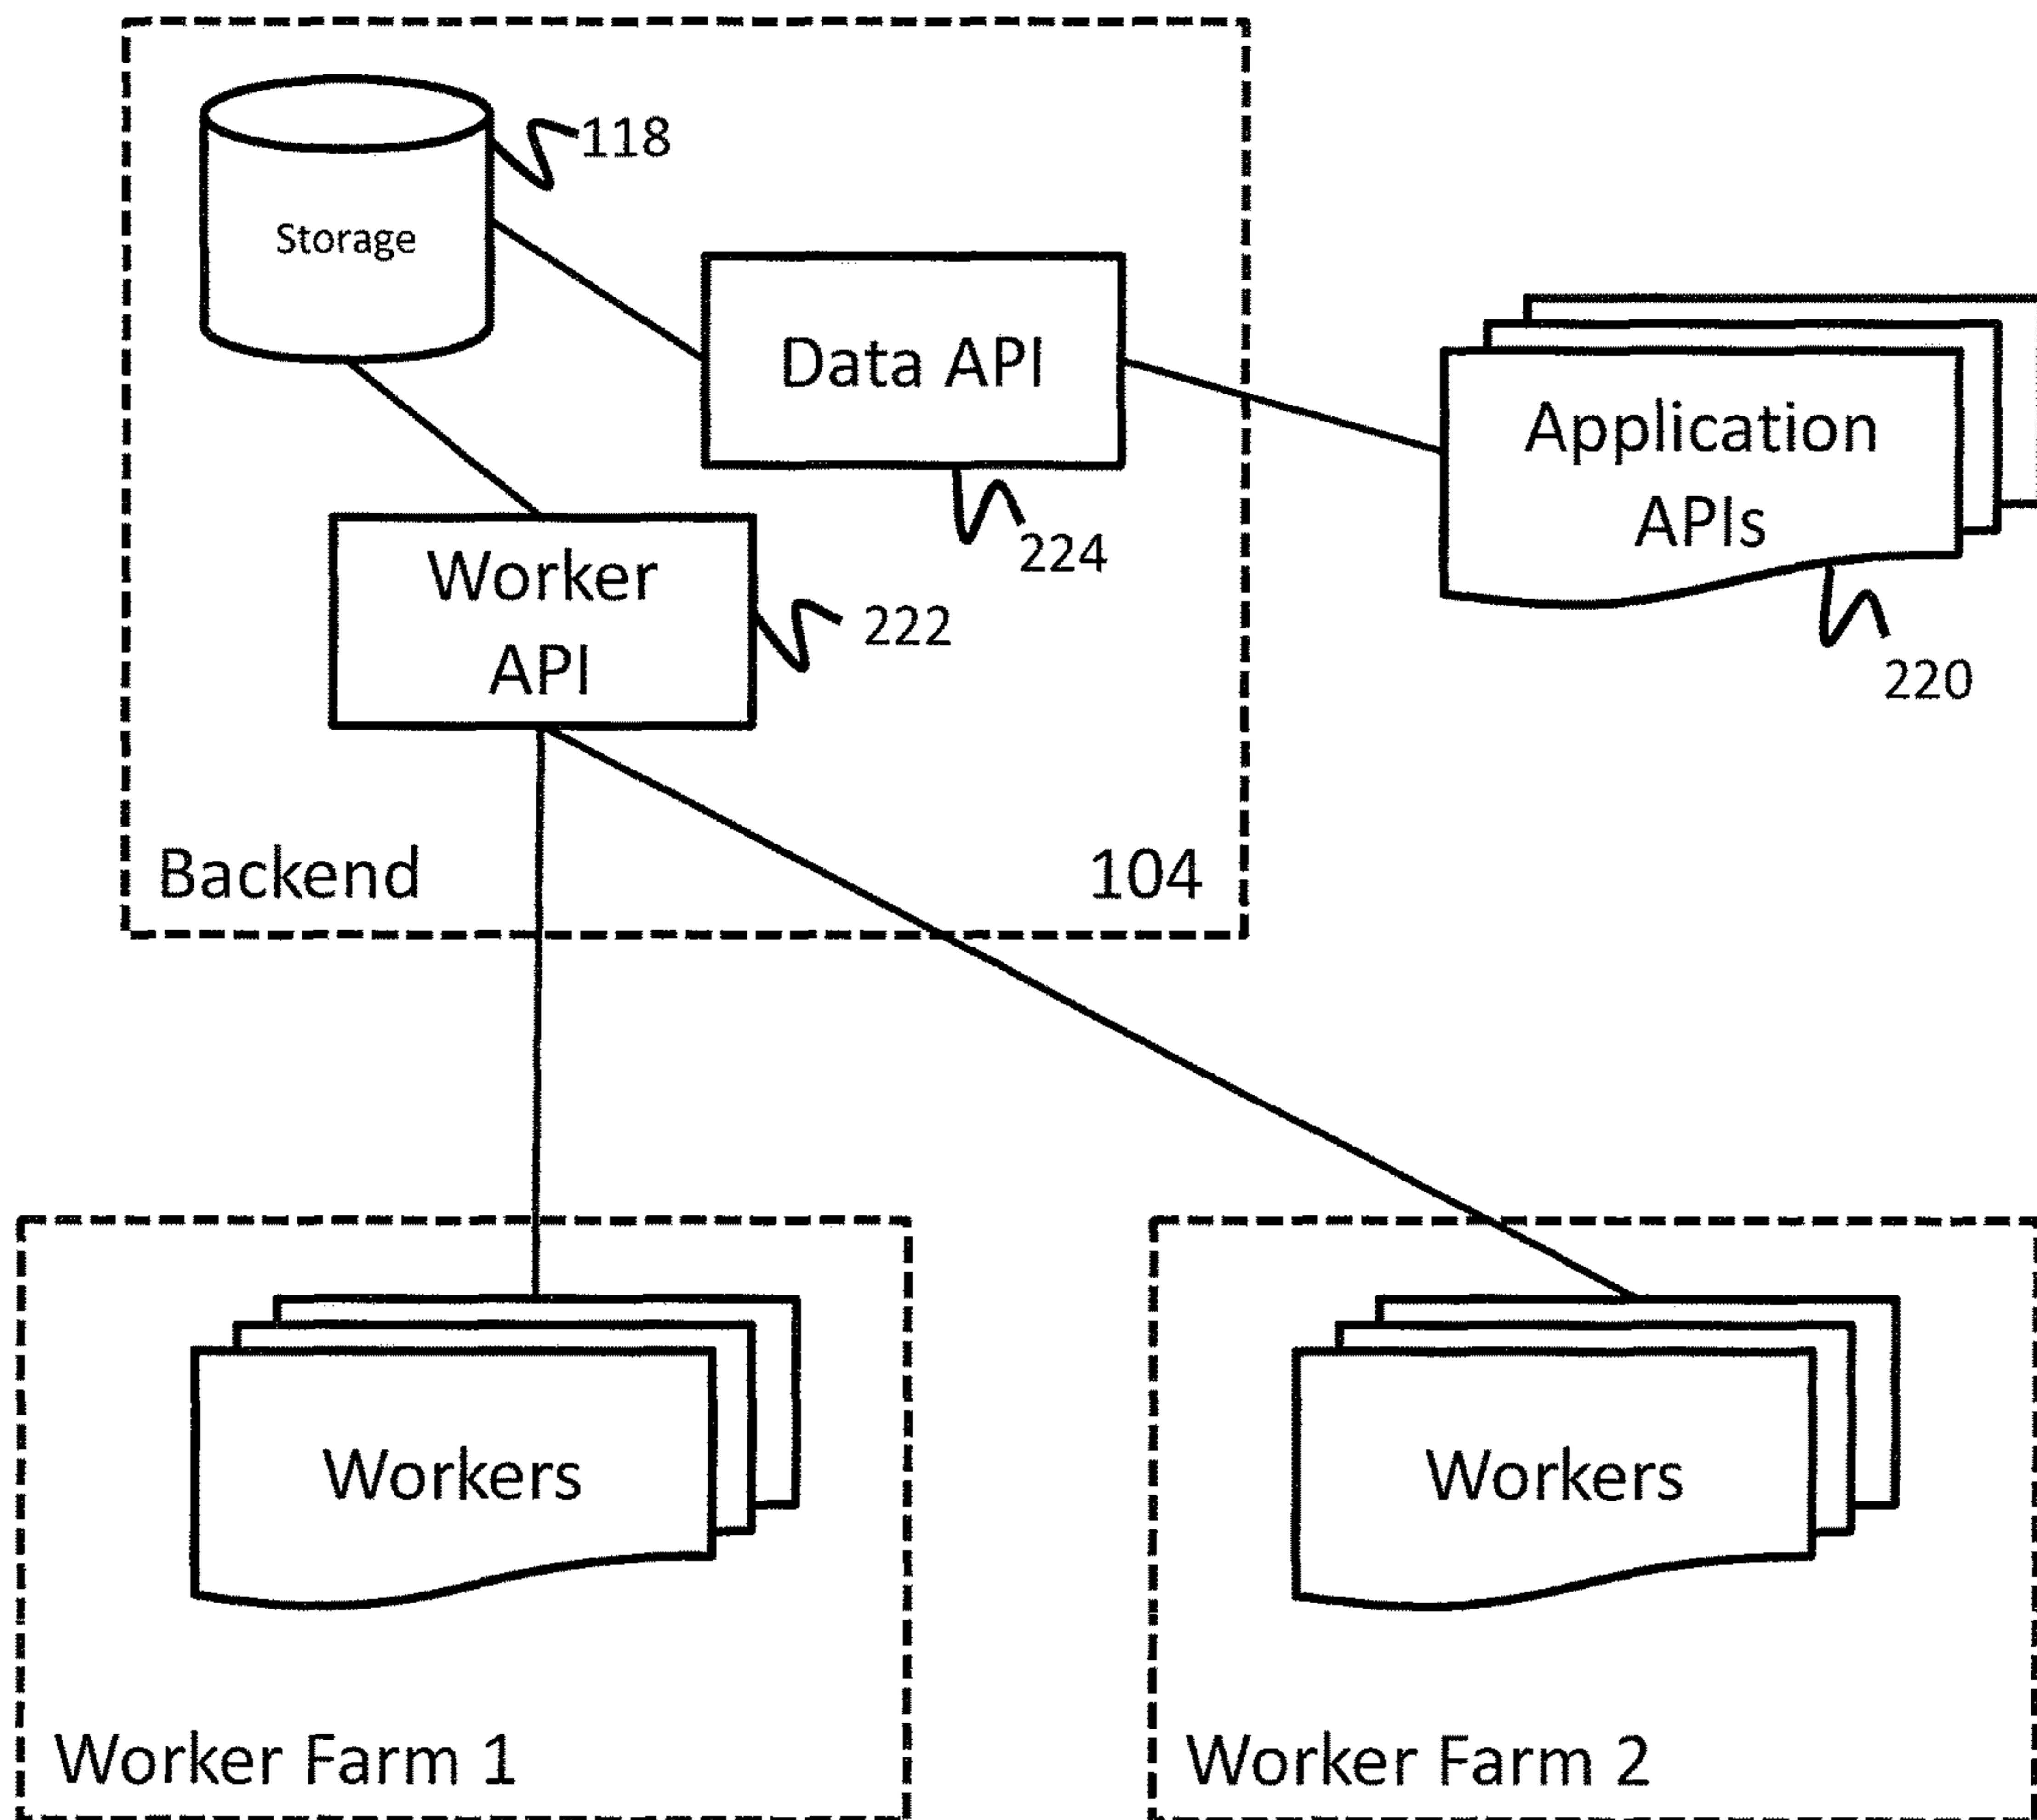


FIG. 1A

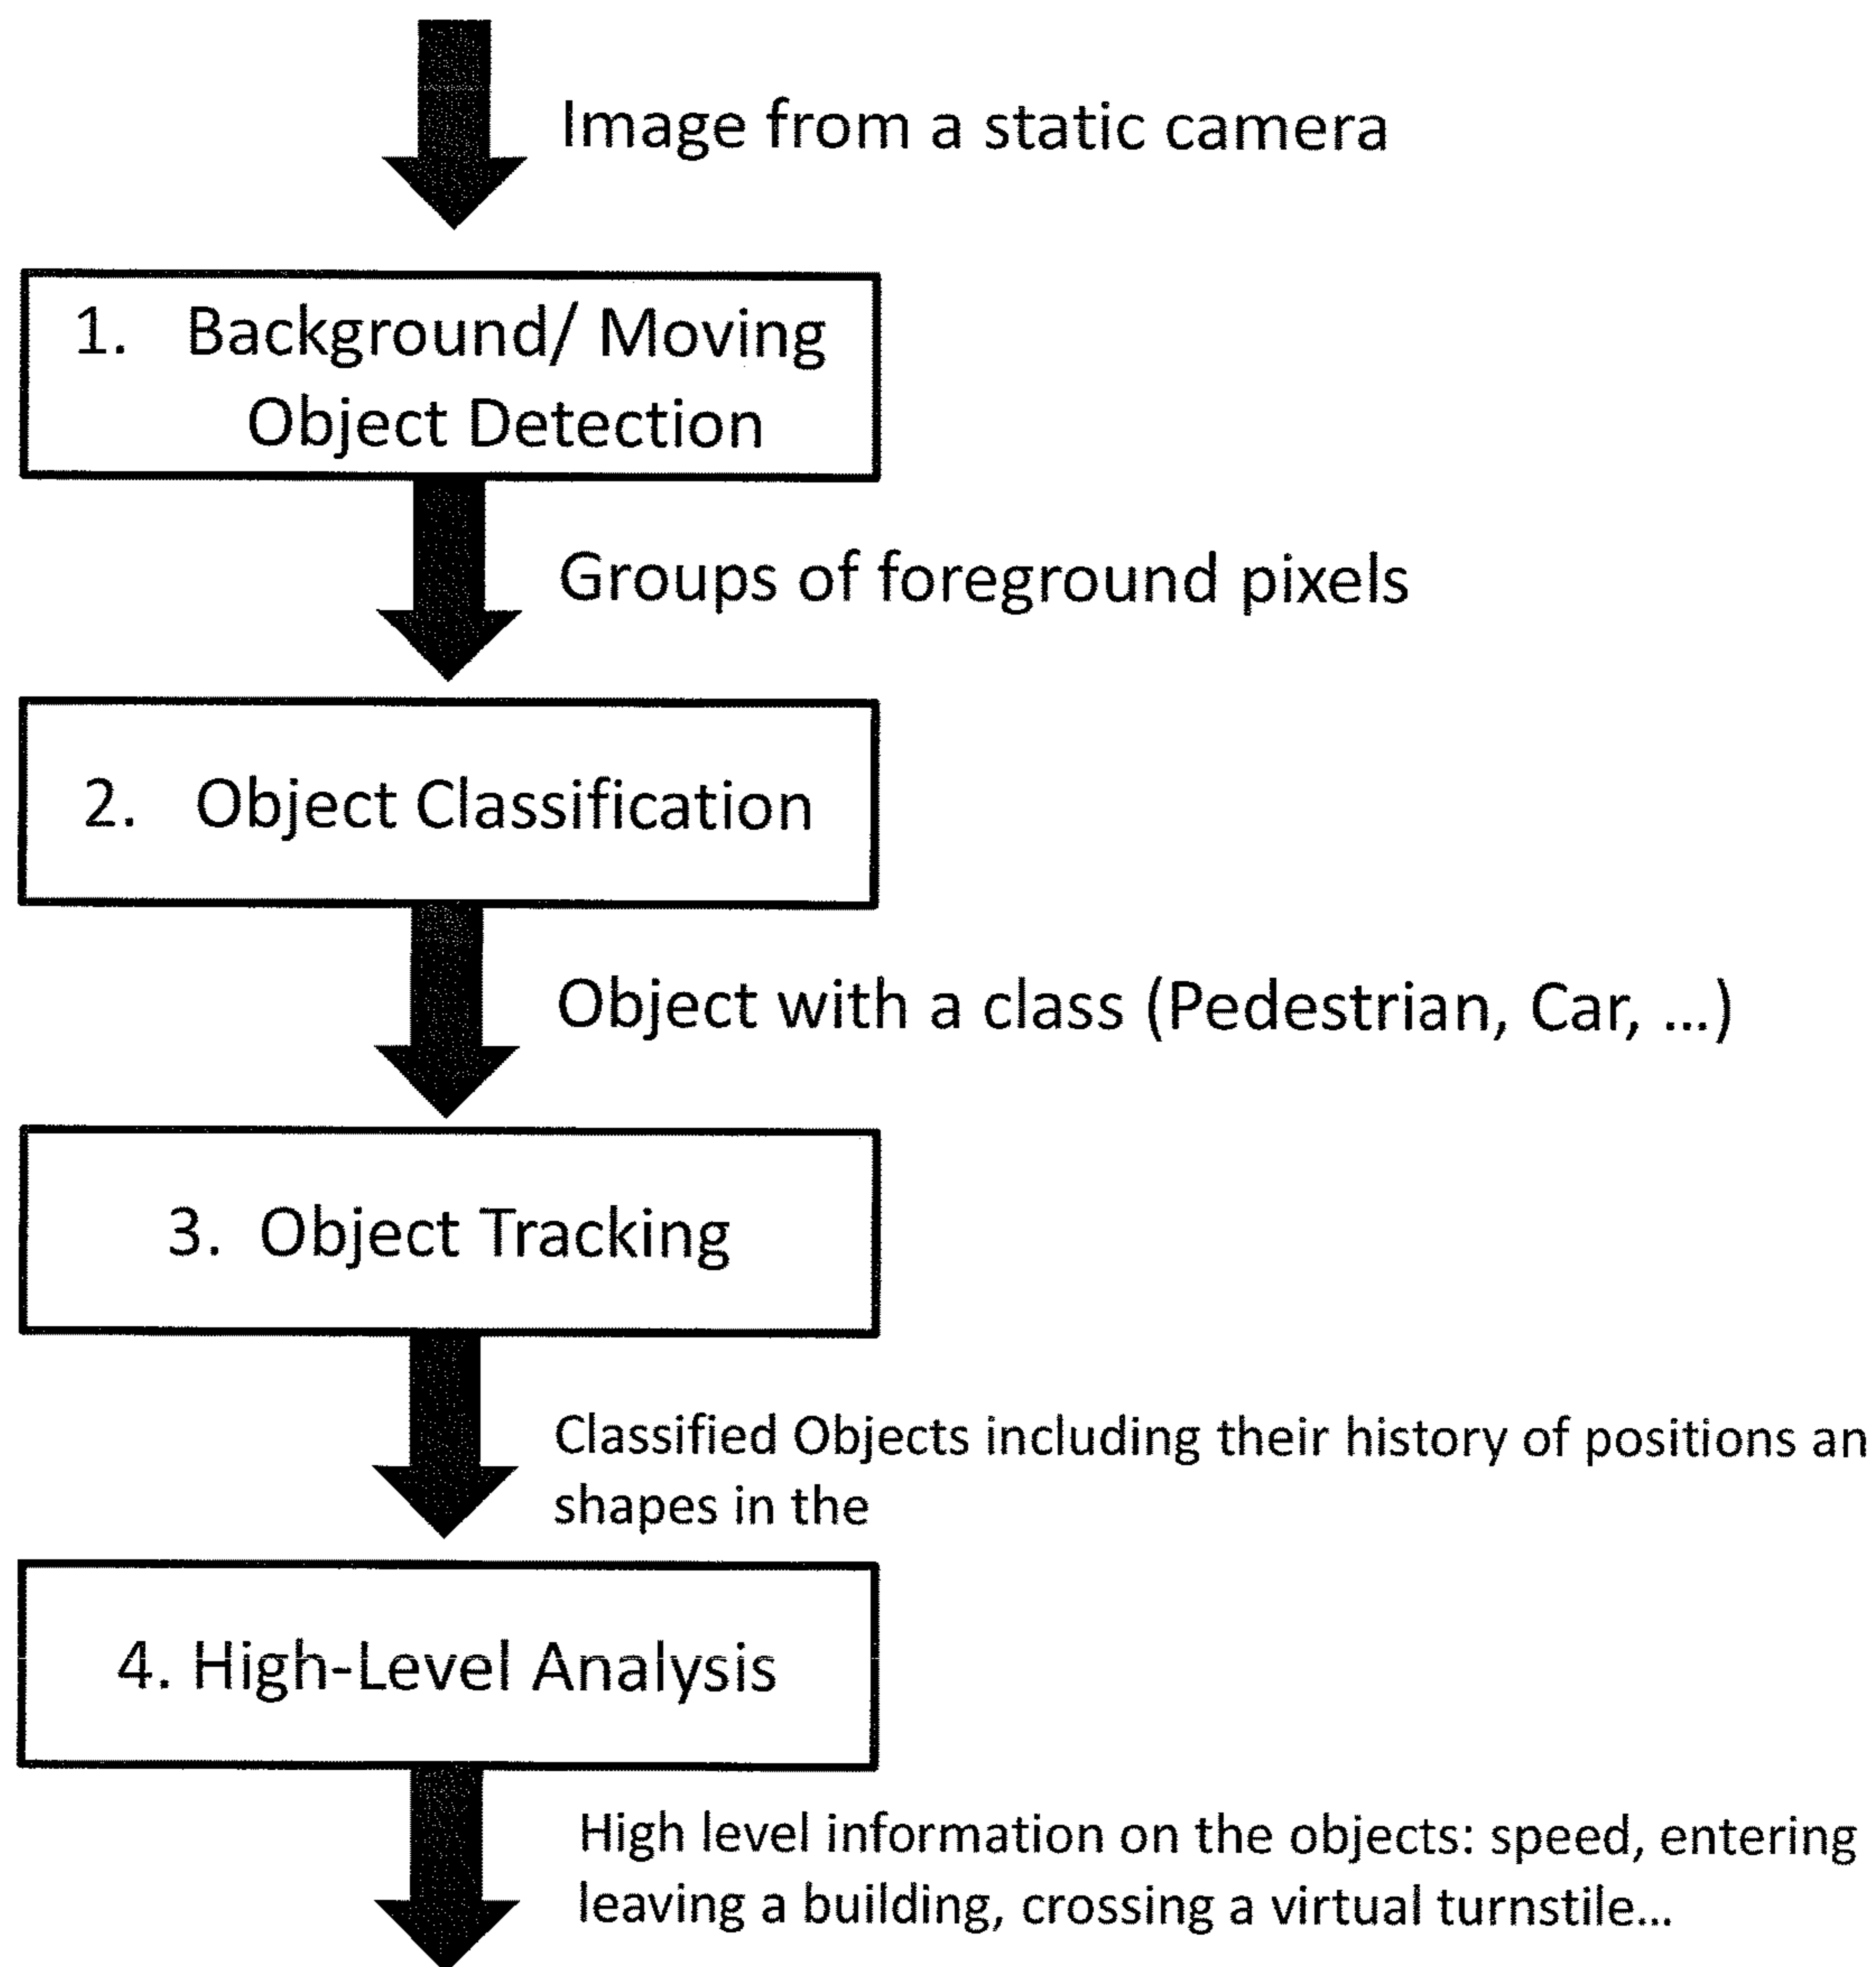


FIG. 2

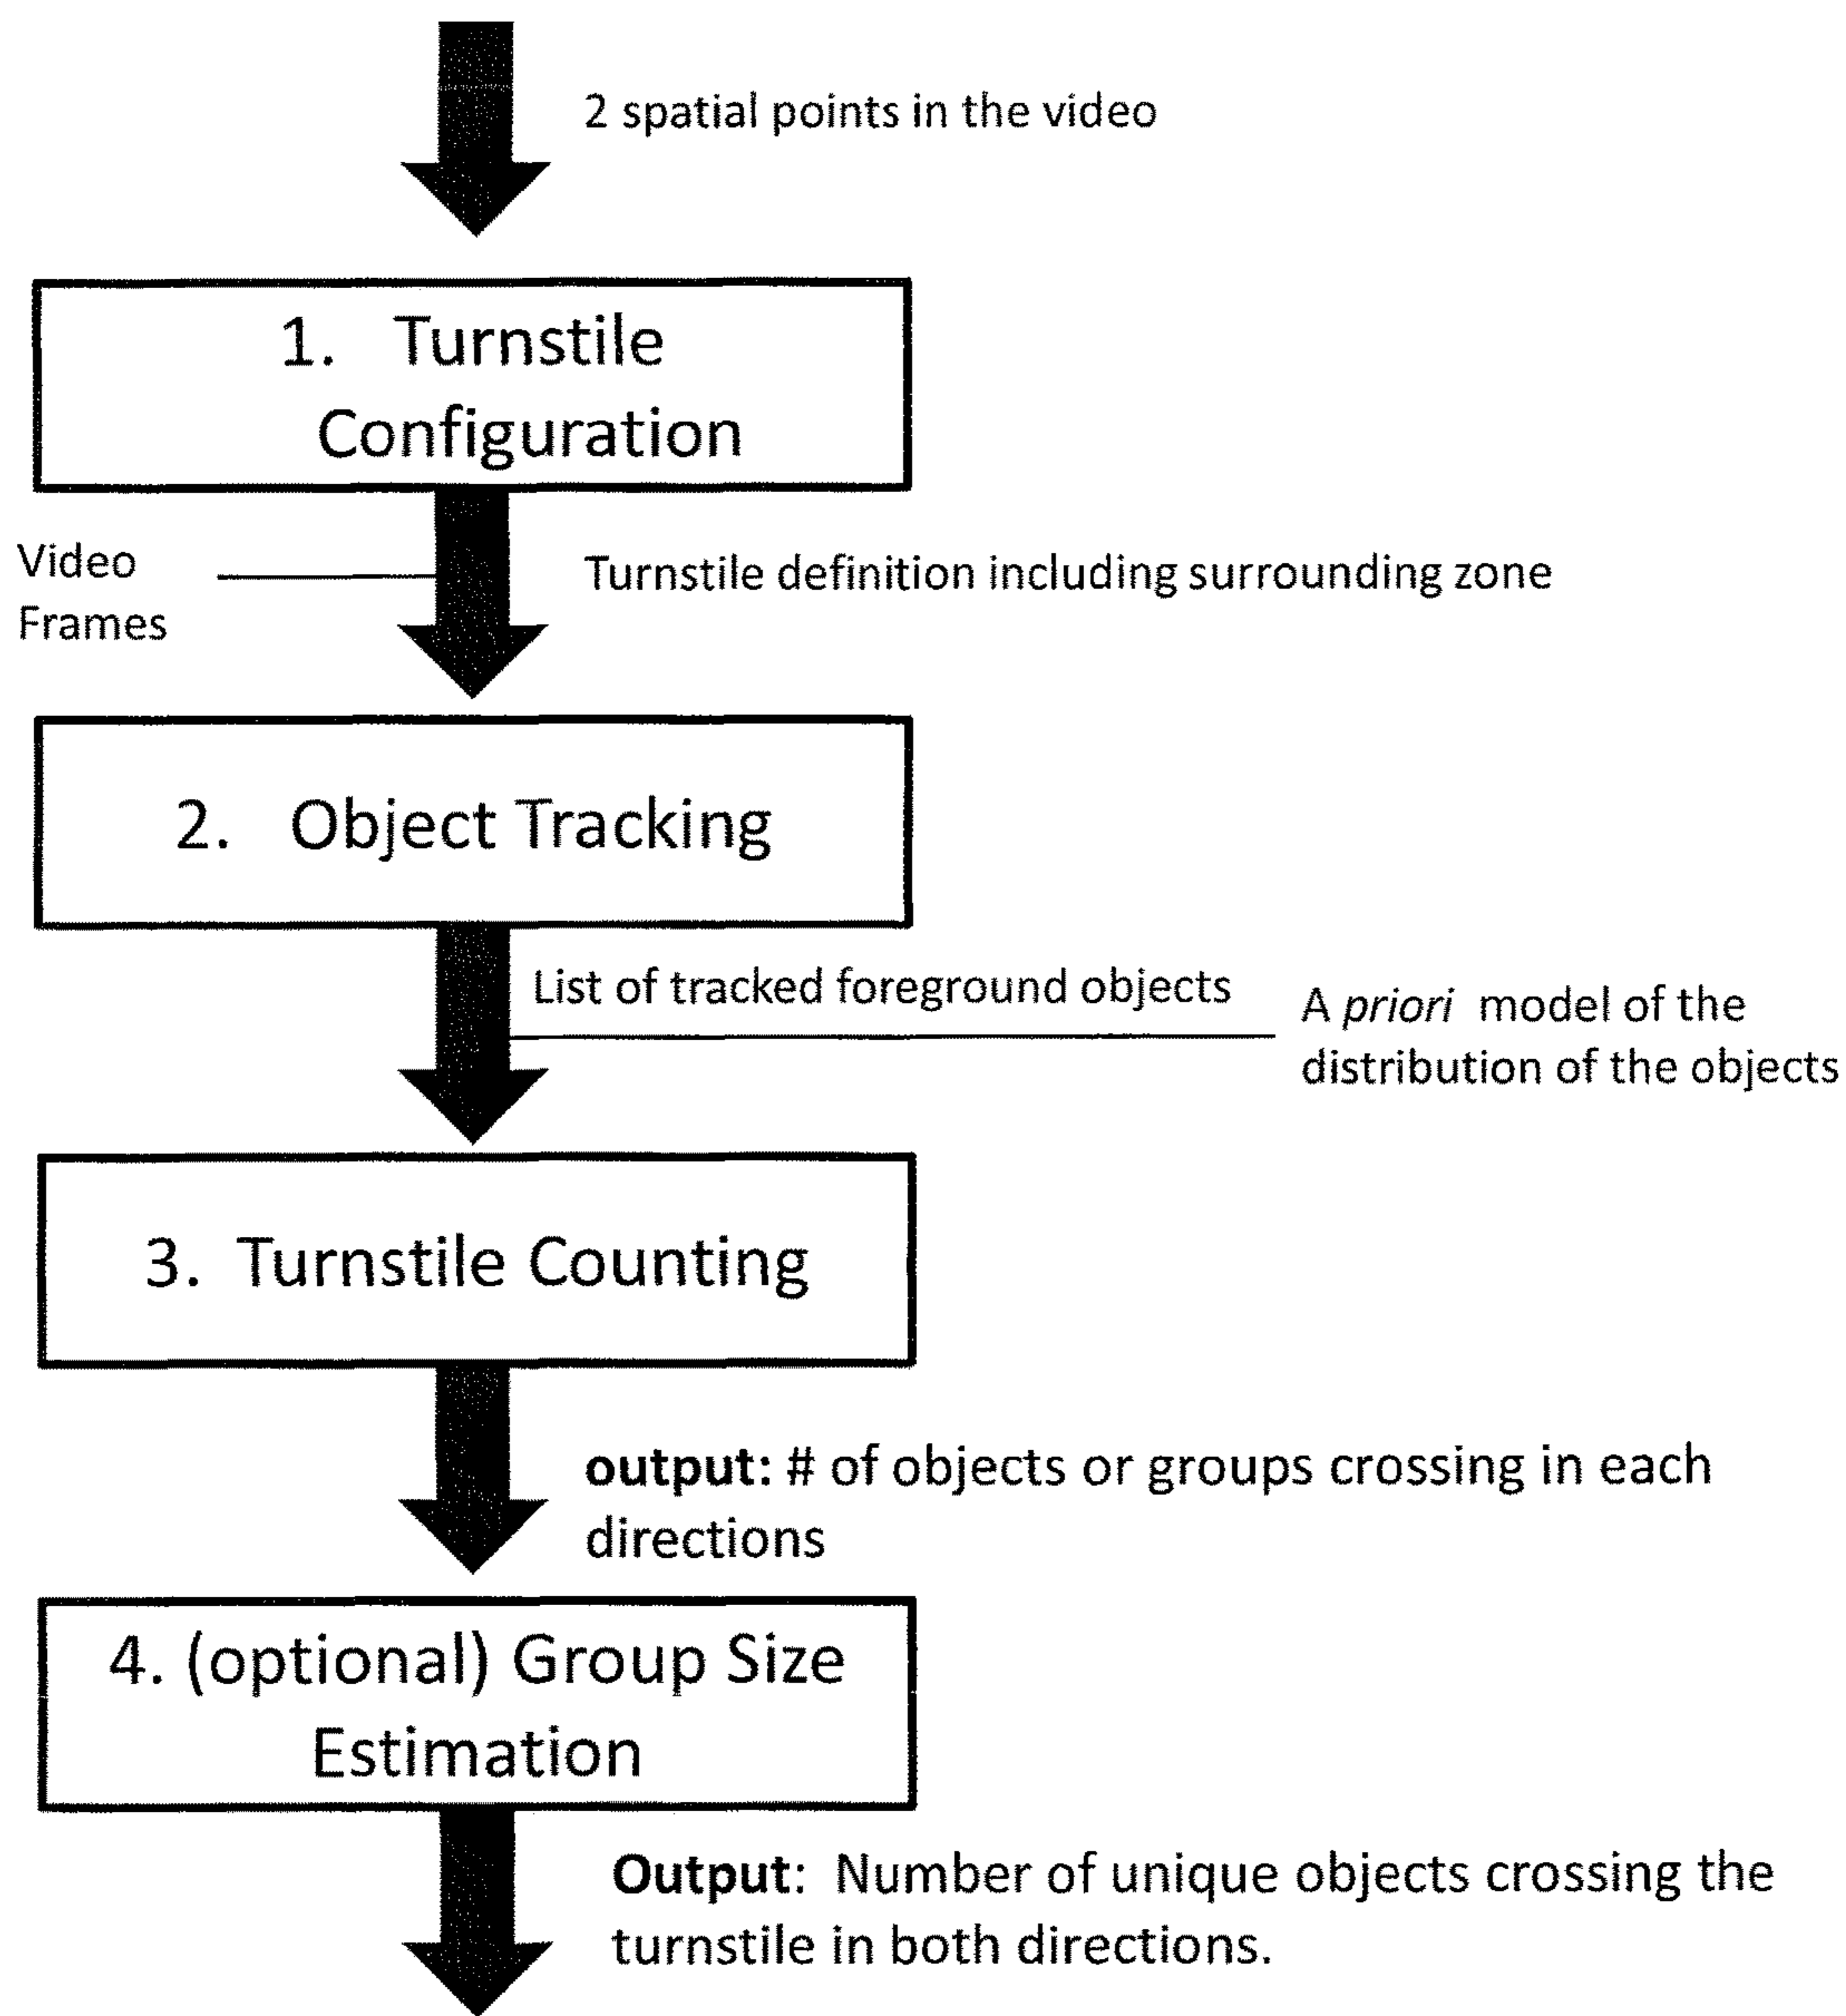


FIG. 3

Virtual Turnstile
illustration

The Virtual Turnstile

9/11
MEMORIAL


do not touch
move this
ment from this
SECTION.

Thank you.

Larry Nunziato
Senior Security Technology Coordinator
World Financial Center
Ph. (212) 312-1010 Cell: (646) 338-0840

86% average
98% in normal

Just one smart phone on a window!



TS#=1 SZ=70 VIS=0 +0/0-0/0

FIG. 4

VIRTUAL TURNSTILE SYSTEM AND METHOD

RELATED APPLICATIONS

This application claims priority from U.S. Provisional Application No. 62/150,646, filed Apr. 21, 2015. This application is also a continuation-in-part of U.S. patent application Ser. No. 14/727,321, filed Jun. 1, 2015, which claims priority from U.S. Provisional Application No. 62/005,815, filed May 30, 2014. All the foregoing patent applications are incorporated by reference herein in their entirety.

FIELD OF THE INVENTION

Embodiments disclosed herein are in the field of video data processing, in particular object detection.

BACKGROUND

Video cameras have been used to capture video images for various applications such as surveillance, machine vision, security monitoring, inspection, sensing and detection. Digital video data is processed using various computer-implemented techniques to extract useful information or to alter the video data itself. The digital video data may be sent to a nearby or remote image processing device to be processed and displayed. As the number of digital image sources increases (surveillance cameras, cell phone cameras, etc.) more capable and efficient techniques for extracting information regarding the world in near real-time are desirable.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a block diagram of a video data processing system according to an embodiment.

FIG. 1A is a block diagram of components of the video data processing system according to an embodiment.

FIG. 2 is a flow diagram of a video data processing method according to an embodiment.

FIG. 3 is a flow diagram of a method for processing output of a video data processing method according to an embodiment.

FIG. 4 is a diagram illustrating a virtual turnstile according to an embodiment.

DETAILED DESCRIPTION

Embodiments described herein include a system and method for video data processing. Video data from multiple streaming sources is processed in order to determine the status of various aspects of environments. The video data processing system uses video streams to measure activity levels in the physical world. This provides information that enables people and businesses to interact more effectively and efficiently with physical locations and cities.

In an embodiment, the video data processing system uses input video streams from a variety of sources. Sources include existing video feeds such as security cameras, video feeds contributed by system users through old smartphones placed in a fixed location simple webcams, or embedded sensors that contain a video feed and some video analysis software. The system includes a backend subsystem consisting of specially programmed processors executing soft-

ware that manages video feeds, processes the video feeds into data, stores the data, and computes analytics and predictions.

Embodiments facilitate the processing and analysis of any possible video source, whatever its type or support. These sources include: existing public video cameras in the form of standard video streams; existing public video feeds in the form of .jpg files regularly updated on a website; simple webcams; security cameras installed for security purposes but whose feeds can be ingested by the video data processing system to extract data; video streams or files coming from old cell phones that run a video sensing application specific to the video data processing system. The sensing application can either produce actual video streams or encore video files, and pushes them on a hosted storage server such as FTP server or Amazon S3 server. Using a smartphone as a video sensor, a capture device is configured to stream data out through files. This solves a major problem of setting cameras and exporting their feeds to a different network location on the internet.

The system thus provides a unified framework to intake video frames coming from these various described sources, and to unify their geolocation and time reference so as to be able to compare any geolocated or time stamped data extracted from them.

In an embodiment using a smart phone, consecutive video files on the smartphone are encoded, time stamped, and pushed to an FTP server to produce a stable stream of video content without having to have a video streaming server in the loop, but rather a simple file server.

These video feeds are produced by multiple types of entities, including: companies or entities that own video feeds, and provide them for free—e.g. the DOT in New York; companies or entities (e.g. retailers) that own video feeds, and provide them to the video data processing system in exchange for having them transformed into valuable data; companies or organizations that are paid for access to the video feeds they own and operate—e.g. earthcam; companies with whom there is no monetary exchange, e.g. they provide their feed, in exchange for a minimal amount of data for free; and individual contributors who use old smart phones or contribute old cell phones which are hung on windows or wall surface. By running the sensing application on these old phones, new video data processing system video feeds are created.

Compiling video data from many different sources to create data insights and analytics has more scaling network effect than all single data sources combined. This is made possible in part by aggregating data from multiple sources (including individual contributors) into a combined, stable source.

Embodiments include various video algorithms dedicated to transforming a video signal into data and measurements. Embodiments further include data algorithms that combine measurements from video feeds with lower resolution activity maps, weather information, and local event data, to infer place activity in space and time. An output interface includes tools to turn the data extracted from videos into human readable information and useful actions.

FIG. 1 is a block diagram of a video data processing system architecture **100** according to an embodiment. Input video sources **102** can include public cameras, security cameras, mobile applications, and any other device capable of capturing digital video streams. This includes a proprietary protocol to do continuous streaming over long periods of time. A feedback loop is used to reconstruct data, and to identify and solve bandwidth issues.

The input video sources **102** are very varied in nature and quality as previously described. A backend subsystem **104** receives video data streams from the input video sources **102**. Feed management module **112** receives the video data streams. Other management modules include a worker management module **114**, a locations management and geographic intelligence module **116**, and a data storage module **118**. As used herein, “worker” implies one or more servers and one or more processors for processing data. Workers can be distributed geographically, and processing tasks may be distributed among workers in any fashion. Data storage module **118** is shown as a single module existing in backend **104**. However, actual data storage can be, and typically is, distributed anywhere over the internet. Data storage module **118** is thus a data management module and possibly actual data storage, but not all data will be stored locally.

Input video sources **102** also communicate with a contributor management module **110**. Contributor management module **110** oversees and tracks the various input video sources, including their locations and “owners”. In some instances, individual owners are paid for making their video data available to the system. Video analysis workers **106** represent multiple special purpose processors tasked with executing video analysis worker processes as further described below. Analyzed video data is stored by data storage manager **118**, and also further analyzed by data analytics module **108** as further described below. Data analytics module **108** represents special purpose processors executing data analytics processes. Data analytics module **108** further has access to external data sources **122**, which provide data such as weather information, event information related to a location of the video data, etc. Data analytics module **108** may combine external data with the output of the video analysis workers **106** to produce more meaningful output data that is stored by data storage management **118** and output to user interface and user applications **120**. User interface and applications **120** make processed video data available to users in a highly accessible form. User interface **120** is available in various embodiments on any computing device with processing capability, communication capability, and display capability, including personal computers and mobile devices.

In an embodiment, backend **104** is a multi-layered system whose roles include: registering all existing video streams and their sources; if the source is a contributor, storing availability and contact information to provide data or to pay them, based on the availability of their sensors; managing “worker” processes that process all video feeds in a different subsystem, and will report data to backend **104**; gathering and storing data extracted from video streams; consolidating and merging all data from various sources (e.g., video measurements, weather APIs); packaging and serving data for applications or as an output of backend **104**; and architecturally removing the dependency of the video algorithm processor on the various sources of data.

According to one aspect of backend **104**, it serves to coordinate the distribution of all input sources and worker processes over different types of networks and environments. FIG. 1A is a block diagram illustrating this concept. A worker application programming interface (API) **222** is defined for worker processes running on worker farms such as worker farm **1** and worker farm **2**. Worker processes are software processes that can run on any operating system such as Linux, Mac OS or Windows with sufficient random access memory (RAM) and central processing unit (CPU) power to register with backend **104** over a Wide Area Network (WAN), assign individual video feeds with corre-

sponding algorithm parameters and specifications, and then start processing that feed and reporting output data back to the same API **220** endpoint family. These worker processes also report on their level of activity, CPU and RAM occupancy, as well as availability. Thus the video data processing system can instantiate worker processes in many varied environments, leveraging available platforms that can contribute CPU power.

Various applications APIs **220** can be used to allow various applications to communicate data to data APIs **224**.

The video data processing system executes various video algorithms and various data algorithms. In an embodiment, the video algorithms are based on a layered stack of algorithms. In an embodiment, these algorithmic layers are based on the assumption that video feeds have a static viewpoint and an average frame rate greater than 0.2 frames per seconds, but embodiments are not so limited.

Video Algorithms

Moving object detection is a layer that detects moving objects or moving parts in the image. It is based on estimating the background image of a fixed video stream, by modeling each point using a Gaussian distribution of values on each channel of a color image, or the amplitude of the combined channels. Each pixel is then modeled as: Gaussian distributions for all channels of the color image; and a Gaussian distribution for the pixel luminance expressed as a linear combination of the three color channels.

Such a model is created and stored in memory for each coordinate point of an image. As new frames arrive in the system, the Gaussian model estimation is updated with the new values of each pixel at the same coordinate by storing the sum S of the pixel values over time, and the sum T of squared values. Given the total number of observations is N , the average and standard deviation of the Gaussian model can then be evaluated as S/N for the mean value and $(2S-S^2)/N$ for the square value of the standard deviation.

In order to adjust the Gaussian values to potential changes in the mean and standard deviation, these values are computed on moving time windows. In order to reduce the complexity of computing all values over a moving averages, a half distance overlapping scheme is used. If M is the minimum window size (number of samples) over which mean and standard deviation is to be estimated, two sets of overlapping sums and square sums are constantly stored: the current sum set and the future sum set. Each set has the number of samples and the sum of values and the sum of square values. When the first set reaches M samples, the second set is reset, and then updated with each new frame. When the first set reaches $M*2$ samples, the future set reaches M samples. The future set values are then copied into the current set values, and the future set is reset. This way, at any point in time after M first samples, the estimation of the Gaussian model always has more than M samples, and it is adjusted over time windows of $M*2$ samples. M is typically set to values ranging from 10 to 1000 depending on applications and video frame rates.

Once a new frame comes in, for each pixel location in an image, it is first assessed whether the current value is part of the background or not. To do so, the normalized distance of the current pixel values is computed for each color channel with the background mean values for each channel. The normalized distance is the distance of the current point to the closest mean adjusted with the standard deviation for the background images. This distance is then normalized towards the amplitude of each channel or the average of all

channels. The raw distance calculated from above is divided by a uniform factor of the average values.

If this normalized distance is greater than a predefined threshold, the pixel is classified as a foreground pixel and assigned to the moving objects. If not, the pixel is deemed as part of the background, it is not assigned to front end masks but is used to update the current background models.

At any point in time, the algorithm assumes that there could be a rapid change in background, so it maintains a candidate background point. That point is either updated or created for each point detected as a foreground point.

If the image is too large, the image can be subsampled by an integer factor to evaluate a lower resolution version of the background. Also, the background statistics can be updated only once every n frames. This is very efficient to make the algorithm “real time” whatever the dimension or frame rate of a video. The CPU occupancy of such a process is controlled and defined with these two parameters. This is a unique way to linearly adjust algorithm reactivity and accuracy based on available or desired computation power.

The object classification layer classifies moving foreground objects (described with reference to the previous layer) into classes of known objects or “noise”. In one embodiment, a customized version of the Haar Pyramid approach is used here. Once all moving objects have been detected, they are classified using a classic supervised learning approach, based on the Haar-like feature Cascade classification (as described in P. A. Viola, M. J. Jones: Robust Real-Time Face Detection. ICCV 2001).

According to embodiments, the system is trained and tested, and the algorithms run only on moving objects, thereby reducing the possibilities and variety of the training and input sets of images. In short the classification scheme only needs to distinguish moving urban objects from each other, as opposed to recognizing one type of object from any other possible matrix of pixels.

A tracking layer detects the trajectory of one given object over time. The system uses a novel approach based on a holistic model of the trajectories in the image based on existing known foreground objects or newly emerged objects.

An analysis layer uses the type and trajectory information to detect higher level, human readable data such as vehicle or pedestrian speed, and people entering or exiting a location. Inferences can also be drawn based on building layouts, vehicle traffic flows, and pedestrian traffic flows.

FIG. 2 is a flow diagram of a process of object detection, classification, object tracking and analysis according to an embodiment. With reference to FIG. 2, embodiments of video algorithms detect and track foreground objects from static cameras and infer high level information such as object count and speed in real-time. They are based on a four-layer stack: object detection, classification, tracking and analysis. These algorithmic layers are based on the assumption that video feeds have a mostly static viewpoint and a frame rate greater than 0.2 frames per second.

A background/moving object detection process (1) takes as input an image from a static camera. Process (1) outputs groups of foreground pixels and a dynamic model of the background.

A goal of this layer is to detect moving objects or moving parts in the image. It is based on estimating the background image of a fixed video stream, by modeling each point using a Gaussian distribution of values on each channel of a color image, or the amplitude of the combined channels. For color images, the value of each channel is modeled as a Gaussian distribution. The Gaussian model, defined by its 0 and 1

moment, or mean and variance, is created and stored in memory for each coordinate point of an image.

In order to determine if a pixel p is part of the background or foreground, we compute a normalized distance metric of a pixel to the background value as the linear combination of the normalized distance for each channel. A normalized distance for a channel is defined as the absolute distance from the pixel channel value to the background value, divided by its standard deviation.

$$d(i)=(v(i)-b(i))/\sigma(i)$$

where (i) is the channel index, v the current pixel channel value, b the background channel value and σ the current estimate of the background channel value standard deviation.

$D=d(0)+d(1)+d(2)$ if the image is a 3 channel color image.

As new frames arrive in the system, we first compute the normalized distance of this pixel to the current background Gaussian model. If the value is less than a first threshold $T1$, we consider the pixel as part of the background and update the Gaussian model for this coordinate point with the current pixel values.

If the value is greater than $T1$, we create a new model for a new background candidate. Things might have changed in the image and we need a new background candidate to adjust to these changes. If there are already some background candidates available, we first compute distances of the current pixels to other candidates. If any distance is less than $T1$, we update the best matching candidates (the one with the lowest distance) with the current value. If no match was found, we create a new candidate.

If a candidate was not updated for a given period of time S , we cancel the background candidate.

Each candidate has a lifetime span, that is equal to the time elapsed between its creation and its last update. The lifetime span cannot be greater than a parameter E called eternity.

$$LS=\text{MIN}(E,t(\text{updated})-t(\text{created})).$$

If any of the candidate backgrounds has a longer lifetime span than the current background, we cancel the current background value and replace it with the new, longer lifetime value. This helps the model adjust to rapid background changes.

If the distance metric is greater than a different factor $T2$, we mark the pixel as being part of the foreground.

In order to adjust the Gaussian values to potential changes in the mean and standard deviation, we estimate all Gaussian model values over overlapping time windows. In order to reduce the complexity of computing all values over moving averages, we use a half-overlapping scheme. If M is the minimum window size (number of samples) over which we want to estimate Gaussian models, we constantly store two sets of overlapping sums and square sums: the current sum set and the future sum set. Each set stores the number of samples and the sum of values and the sum of square values that help compute mean and variance. When the first set reaches M samples, we reset the second set and start updating it with each new frame. When the first set reaches $M*2$ samples, the future set reaches M samples. We then copy the future set values into the current set values, and reset the future set. This way, at any point in time after M first samples, we always have an estimation of the Gaussian model that has more than M samples, and adjust over time windows of $M*2$ samples. M is typically set to values ranging from 10 to 1000 depending on applications and video frame rates. As a result, outside of the starting period

where we have less than M samples processed in total, all our Gaussian model estimates rely on at least M samples.

In order to reduce computation cost, we can subsample the image spatial reference by a factor P. Also, we can subsample the time reference by another factor Q—we update the background statistics only once every Q frames. This reduces the number of operations needed significantly. However, the foreground estimation cannot be subsampled, so complexity is only reduced for background estimation. This is very efficient to make the algorithm real time whatever the dimension or frame rate of a video. We actually control and define the CPU occupancy of a such a process with these two parameters. This is a unique way to linearly adjust algorithm reactivity and accuracy based on available or desired computation power.

An object pixel classification process (2) takes as input groups of foreground pixels. The output is one or more objects per group with an associated class.

The goal of this layer is to classify foreground described from process (1) above into classes of known objects or “noise”. In an embodiment, a customized version of the ADABOOST Cascade approach is used.

Once we have detected all moving objects we classify them using a classic supervised learning approach, based on the ADABOOST Cascade classification (described in Viola and Jones P. A. Viola, M., and J. Jones: Robust Real-Time Face Detection. ICCV 2001).

Embodiments of the method train, test and run the algorithm only on moving objects, thereby reducing the possibilities and variety of the training and input sets of images. In short our classification scheme only needs to recognize moving urban objects from each other, as opposed to recognizing one type of object from any other possible matrix of pixels.

This step also helps separate groups or aggregates in some cases—a car and pedestrians are close to each other and detected as the same object, we will be able to detect them separately in many occasions, thus splitting the original object in two separate objects.

An object tracking process (3) takes as input an instance of one object at one point in time, with or without associated class. The output is a linked appearance of the same objects at different times, with trajectory and shape over time.

The goal of this layer is to connect the occurrence of the same object in consecutive frames so as to understand the object’s movement in the image.

At each new frame, we try to match new foreground objects with existing, connected objects tracked in prior iterations, or if no match is found, we create a new object. We use a combination of shape, predicted position based on previous motion, and pixel content, to do the matching.

An object trajectory analysis and classification process (4) takes as input objects with trajectories, and outputs high level information on objects.

The goal of this layer is to use the type and trajectory information to detect higher level, human readable data such as vehicle or pedestrian speed, and people entering or exiting a location. We can also infer on the building layouts based on traffic flows of pedestrians and vehicles.

Data Algorithm: Virtual Turnstile

FIG. 3 is a flow diagram of a virtual turnstile method according to an embodiment. The virtual turnstile can be visualized as a fictive line segment materialized by two spatial points in a video.

The virtual turnstile counter is an algorithm that counts the number of objects (people, cars, etc.) crossing a virtual line drawn on a video, possibly breaking down these counts by object classes (color, kind, male, female, young, old . . .) or any other information related to these objects.

Turnstile configuration (1) is defining the line where we want to count people, that is two points materializing this line in the video—manually entered or automatically deduced by the Scene Analysis algorithm. We limit the computation and detection of moving objects to a mask surrounding the turnstile, and as wide as the turnstile width. We also define a main crossing angle, which is the average angle objects are expected to cross the turnstile with. Then we define the maximum tolerated deviation to this angle: if an object crosses the turnstile with an angle that is different from the main angle to an extent larger than the maximum deviation, the object will be discarded.

Object detection and tracking (2) is as described with reference to FIG. 2.

The turnstile definition, including the surrounding zone, is input to the object tracking module (2). Output of the module is input to the turnstile counting or trajectory analysis module (3). This includes a count of objects crossing the virtual turnstile in the right conditions. This layer computes the crossing angle of the object to the turnstile and compares it to the angles definitions of the turnstile. In order to remove noise, this layer also removes any object whose trajectory is too erratic, e.g., where the virtual force to change trajectory is too large. Starting from Newton’s equation:

$$(\text{sum of forces})=(\text{acceleration})\times(\text{mass of object})$$

Mass as a function of object surface is estimated. If the norm of the acceleration multiplied by the average surface of an object over time is higher than a threshold T, we discard the object. This threshold is estimated by computing all forces over some time over one image, and selecting the 99 percentile of this distribution.

Group size estimation (4) is optional. A group of N people walking together could be detected as one single blob, and in the algorithm presented so far, would simply be counted as one, and not N. One approach is to compute the median size of a person, and then use that to estimate the number of persons present in a blob. In an embodiment, the algorithm self-calibrates in a first step, assessing the average dimension of a person. To do this it registers all the sizes of moving objects crossing the threshold. It assumes they are mostly people. It also assumes that at least a given percentage N of them are not merged into groups. If we order all sizes, the $N*\text{tot_elements}/100$ element represents the typical person size. We use this to infer how many people are in a blob that crosses the line. We assume that there is always a certain overlap between people and factor that in the regression function to go from number of pixels in the blob to # persons in the blob. Another approach is to use a ground truth approach to adjust parameters and have a linear per part regression function: we design a liner-per-part model (a set of connected segments) as the function $f(x)=n$ where x is the # pixels and n the estimated number of people. We use a linear regression to start, and then split the set in two subsets where the samples are divided by two, and iterate until reaching the desired number of segments of typically 8. In low quality video footage, it is sometimes impossible to distinguish individual people, so tracking waiting persons to know the time that it takes to travel the whole line is usually not viable. Instead, this approach to estimating the average wait time consists in dividing the line length by the average line speed. The line speed is estimated by computing the

optical flow of a set of salient points into the line contours over each pair of consecutive frames.

This methodology was tested on views of the 9-11 Memorial Park in NYC captured by one smart phone camera, as shown in FIG. 4. Accuracies up to 98% were under normal conditions. Under average conditions accuracy reaches 86%.

Aspects of the systems and methods described herein may be implemented as functionality programmed into any of a variety of circuitry, including programmable logic devices (PLDs), such as field programmable gate arrays (FPGAs), programmable array logic (PAL) devices, electrically programmable logic and memory devices and standard cell-based devices, as well as application specific integrated circuits (ASICs). Some other possibilities for implementing aspects of the system include: microcontrollers with memory (such as electronically erasable programmable read only memory (EEPROM)), embedded microprocessors, firmware, software, etc. Furthermore, aspects of the system may be embodied in microprocessors having software-based circuit emulation, discrete logic (sequential and combinatorial), custom devices, fuzzy (neural) logic, quantum devices, and hybrids of any of the above device types. Of course the underlying device technologies may be provided in a variety of component types, e.g., metal-oxide semiconductor field-effect transistor (MOSFET) technologies like complementary metal-oxide semiconductor (CMOS), bipolar technologies like emitter-coupled logic (ECL), polymer technologies (e.g., silicon-conjugated polymer and metal-conjugated polymer-metal structures), mixed analog and digital, etc.

It should be noted that the various functions or processes disclosed herein may be described as data and/or instructions embodied in various computer-readable media, in terms of their behavioral, register transfer, logic component, transistor, layout geometries, and/or other characteristics. Computer-readable media in which such formatted data and/or instructions may be embodied include, but are not limited to, non-volatile storage media in various forms (e.g., optical, magnetic or semiconductor storage media) and carrier waves that may be used to transfer such formatted data and/or instructions through wireless, optical, or wired signaling media or any combination thereof. Examples of transfers of such formatted data and/or instructions by carrier waves include, but are not limited to, transfers (uploads, downloads, e-mail, etc.) over the internet and/or other computer networks via one or more data transfer protocols (e.g., HTTP, FTP, SMTP, etc.). When received within a computer system via one or more computer-readable media, such data and/or instruction-based expressions of components and/or processes under the system described may be processed by a processing entity (e.g., one or more processors) within the computer system in conjunction with execution of one or more other computer programs.

Unless the context clearly requires otherwise, throughout the description and the claims, the words “comprise,” “comprising,” and the like are to be construed in an inclusive sense as opposed to an exclusive or exhaustive sense; that is to say, in a sense of “including, but not limited to.” Words using the singular or plural number also include the plural or singular number respectively. Additionally, the words “herein,” “hereunder,” “above,” “below,” and words of similar import refer to this application as a whole and not to any particular portions of this application. When the word “or” is used in reference to a list of two or more items, that word covers all of the following interpretations of the word: any of the items in the list, all of the items in the list and any combination of the items in the list.

The above description of illustrated embodiments of the systems and methods is not intended to be exhaustive or to limit the systems and methods to the precise forms disclosed. While specific embodiments of, and examples for, the systems components and methods are described herein for illustrative purposes, various equivalent modifications are possible within the scope of the systems, components and methods, as those skilled in the relevant art will recognize. The teachings of the systems and methods provided herein can be applied to other processing systems and methods, not only for the systems and methods described above.

The elements and acts of the various embodiments described above can be combined to provide further embodiments. These and other changes can be made to the systems and methods in light of the above detailed description.

In general, in the following claims, the terms used should not be construed to limit the systems and methods to the specific embodiments disclosed in the specification and the claims, but should be construed to include all processing systems that operate under the claims. Accordingly, the systems and methods are not limited by the disclosure, but instead the scope of the systems and methods is to be determined entirely by the claims.

While certain aspects of the systems and methods are presented below in certain claim forms, the inventors contemplate the various aspects of the systems and methods in any number of claim forms. For example, while only one aspect of the systems and methods may be recited as embodied in machine-readable medium, other aspects may likewise be embodied in machine-readable medium. Accordingly, the inventors reserve the right to add additional claims after filing the application to pursue such additional claim forms for other aspects of the systems and methods.

What is claimed is:

1. A method for activity monitoring using video data, the method comprising:

a special purpose processor collecting video data from a plurality of dissimilar sources;

the special purpose processor, in an iterative process, defining known foreground objects and newly emerged objects;

the special purpose processor determining moving objects from the video data;

the special purpose processor classifying the moving objects from the video data;

the special purpose processor tracking the moving objects from the video data wherein tracking comprises receiving as input an instance of one object at one point in time and calculating its trajectory and shape over time, and for each new video frame attempting to match foreground objects with existing objects tracked in prior iterations of the iterative process;

the special purpose processor processing the video data from the plurality of dissimilar sources to allow the video data from the plurality of dissimilar sources to be analyzed without dependence on a source of any of the video data; and

the special purpose processor analyzing the video data to generate a count of a number of objects crossing a predetermined line in a time period.

2. The method of claim 1, wherein the video data comprises a plurality of images, and wherein the image data is subsampled by an integer factor to facilitate processing speed.

3. The method of claim 1, wherein the video data comprises a plurality of images, and wherein determining mov-

11

ing objects comprises determining background statistics, wherein background statistics can be updated at variable intervals to affect different processing speeds.

4. The method of claim 1, wherein determining moving objects from the video data comprises estimating the background image of a fixed video stream, comprising modeling points in the video data using a Gaussian distribution of values on each channel of a color image.

5. The method of claim 1, wherein classifying moving objects comprises receiving groups of foreground pixels and outputting one or more objects per group with an associated class using class supervised learning.

6. The method of claim 1, wherein tracking the moving objects comprises receiving an instance of an object at a point in time and outputting a linked appearance of the same object at different times, with trajectory and shape over time.

7. The method of claim 1, wherein analyzing the video data comprises receiving objects with trajectories and outputting information regarding the objects, comprising speed, of objects, objects entering or leaving a location, and inference of building layouts based on movement of objects.

8. A non-transient computer readable medium having stored thereon instructions, that when executed by a special purpose processor cause a method for activity monitoring using video data to be performed, the method comprising:

the special purpose processor collecting video data from a plurality of dissimilar sources;

the special purpose processor, in an iterative process, defining known foreground objects and newly emerged objects;

the special purpose processor determining moving objects from the video data;

the special purpose processor classifying the moving objects from the video data;

the special purpose processor tracking the moving objects from the video data, wherein tracking comprises receiving as input an instance of one object at one point in time and calculating its trajectory and shape over time, and for each new video frame attempting to match foreground objects with existing objects tracked in prior iterations of the iterative process;

the special purpose processor processing the video data from the plurality of dissimilar sources to allow the video data from the plurality of dissimilar sources to be analyzed without dependence on a source of any of the video data; and

the special purpose processor analyzing the video data to generate a count of a number of objects crossing a predetermined line in a time period.

9. The non-transient computer readable medium of claim 8, wherein the video data comprises a plurality of images, and wherein the image data is subsampled by an integer factor to facilitate processing speed.

10. The non-transient computer readable medium of claim 8, wherein the video data comprises a plurality of images, and wherein determining moving objects comprises determining background statistics, wherein background statistics can be updated at variable intervals to affect different processing speeds.

11. The non-transient computer readable medium of claim 8, wherein determining moving objects from the video data comprises estimating the background image of a fixed video stream, comprising modeling points in the video data using a Gaussian distribution of values on each channel of a color image.

12. The non-transient computer readable medium of claim 8, wherein classifying moving objects comprises receiving

12

groups of foreground pixels and outputting one or more objects per group with an associated class using class supervised learning.

13. The non-transient computer readable medium of claim 8, wherein tracking the moving objects comprises receiving an instance of an object at a point in time and outputting a linked appearance of the same object at different times, with trajectory and shape over time.

14. The non-transient computer readable medium of claim 8, wherein analyzing the video data comprises receiving objects with trajectories and outputting information regarding the objects, comprising speed, of objects, objects entering or leaving a location, and inference of building layouts based on movement of objects.

15. A system for processing multiple video data streams, comprising:

a plurality of input video data sources;

a backend subsystem configured to receive video data from the plurality of input video data sources, the backend subsystem comprising,

a plurality of video analysis workers comprising special purpose processors tasked with executing video analysis worker processes, comprising,

in an iterative process, defining known foreground objects and newly emerged objects;

determining moving objects from the video data;

classifying the moving objects from the video data;

tracking the moving objects from the video data, wherein tracking comprises receiving as input an instance of one object at one point in time and calculating its trajectory and shape over time, and for each new video frame attempting to match foreground objects with existing objects tracked in prior iterations of the iterative process;

wherein tracking the moving objects comprises tracking moving objects from a plurality of dissimilar sources to allow the video data from the plurality of dissimilar sources to be analyzed without dependence on a source of any of the video data; and

a data analytics module configured to receive processed video data from the plurality of video analysis workers and output human readable information, wherein the human readable information comprises virtual turnstile information.

16. The system of claim 15, wherein the data analytics module is further configured to access external data sources, including weather data sources, and event information sources.

17. The system of claim 15, further comprising at least one user interface configured to provide the human readable information to a user any computing device with processing capability, communication capability, and display capability, including personal computers and mobile devices.

18. The system of claim 15, wherein the special purpose processors determine moving objects from the video data, and wherein determining comprises:

the special purpose processors classifying the moving objects from the video data; and

the special purpose processors tracking the moving objects from the video data.